

DTIC FILE COPY


2

Robust Planning and Control Using Neural Networks

ONR/DARPA Grant N00014-89-J-3100

Summary of Technical Progress

December 31, 1989



W. Thomas Miller, III
Filson H. Glanz
L. Gordon Kraft, III
Michael J. Carter

Robotics Laboratory
Intelligent Structures Group
Department of Electrical and Computer Engineering
The University of New Hampshire

AD-A217 216

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Contact: W. Thomas Miller, III
Department of Electrical and Computer Engineering
Kingsbury Hall
The University of New Hampshire
Durham, New Hampshire 03824
(603) 862-1326
t_miller@unhh.bitnet

90 01 29 077

Introduction

During the past five years, the Robotics Laboratory of the Department of Electrical and Computer Engineering at the University of New Hampshire has been studying the application of locally generalizing neural networks to difficult problems in control. In a series of theoretical and real time experimental studies, learning control approaches have been shown to be effective for controlling the dynamics of multidimensional, nonlinear robotic systems during repetitive and nonrepetitive operations. This project involves the extension of our work in learning control, with the combined goals of expanding our theoretical understanding of neural network based learning control systems and of extending our experimental work to include hierarchical learning control structures. Our work will consider the efficacy of locally generalizing versus globally generalizing neural network architectures in control applications, as well as developing and analyzing learning control paradigms which are not restricted to specific network architectures. Various robotic systems within the laboratory will form the basis for the real time experimental portions of the research. The concepts explored, however, will be applicable to a wide variety of control problems in addition to robotics. (" ")

In accordance with the above project goals, the ongoing work consists of five parallel efforts: system modeling, task planning, reinforcement learning, control system analysis, and fault tolerance. The project was approved for funding August 1, 1989, and work began September 1, 1989. This progress report summarizes activities during the quarter ending December 31, 1989.

A collection of recent publications has been included as an appendix. Most of these papers represent work that was performed before the start of this grant, but provide background for the efforts currently in progress.

System Modeling

Every control system incorporates some form of model of the system being controlled. Neural network learning appears to be well suited to model building in control since there is often a wealth of training data available from the feedback sensors. The following efforts were proposed to extend our past work in neural network models for control:

1. Investigate neural network architectures which provide rapid performance convergence and which are resistant to learning interference during continuous on-line training in a control system.
2. Investigate non-recursive and recursive neural network architectures for modeling dynamical systems, including the effects of history dependence and time delays.

Research efforts during the current period have focused on the first of these two items. In particular, we have been examining alternative formulations of the CMAC neural network. The traditional Albus CMAC network [Albus, 1975; 1979] utilizes local receptive fields which are rectangular in shape and are distributed along hyper-diagonals in the network input hyperspace. Such locally generalizing networks have been shown to be well suited to learning nonlinear control transformations [Miller, 1986; 1987; 1989; Miller et al. 1987; 1988; 1989], providing rapid and stable training convergence.

Since the receptive fields utilize rectangular window functions (on or off states), CMAC networks build piecewise constant approximations to the learned function. This is sufficient for many applications since the size of the locally constant regions can be controlled as a design parameter. However, some control problems are sensitive to discontinuities in the control transformation. Also, new techniques in neural network based control and reinforcement learning are being developed which utilize the backwards derivative of trained forward models as a local approximation of the system inverse. Determining such backwards derivatives requires that the network function be continuous. Prior work with these techniques at other labs [Jordan, 1988; Werbos, 1989] has involved using traditional multi-layer networks with derivatives computed using an extension of the backpropagation technique. This has been useful for demonstrating the concepts in simulation, but given the limited scalability of such networks, has not been extended to work with real systems.

With these motivations, we have been investigating CMAC neural networks with tapered, rather than rectangular, receptive fields. The general idea of overlapping, tapered receptive fields is certainly not new and



A-1

per US

has a strong basis in biological systems. Albus's original CMAC papers recognized this fact, but utilized rectangular fields as an implementation efficiency. Typical approaches to formulating tapered receptive fields (polynomial splines, radial basis functions, etc.) become undesirably complex as the dimension of the input state space increases. Our work has focused on investigating shapes for tapered receptive fields which are applicable to many dimensional input vectors. Preliminary results indicate that it is possible to create networks which retain the efficiency of the Albus CMAC, but produce piecewise planar functions, rather than piecewise constant functions. In addition to producing continuous outputs and having well defined (and easily obtained) backwards derivatives, such networks tend to converge faster than the traditional CMAC in many problems. However, we have found that the performance of such networks is sensitive to the placement of the receptive fields in the state space (unlike the Albus CMAC), and that the hyper-diagonal arrangement of the Albus CMAC is particularly bad when using tapered receptive fields. We have developed a heuristic procedure for defining the receptive field placement in an N dimensional input space in order to get good performance. We are currently pursuing both extensive experimental validation and theoretical justification for that heuristic. This work has not yet been published. It should be noted that while these new networks deviate considerably from Albus's original formulation in both organization and implementation, we retain the name CMAC to apply to all such locally generalizing networks which are derived from the basic principals proposed in the original work of Albus.

Task Planning

Our past work in neural network based learning control has emphasized the path following aspect of control, which involves successfully carrying out predetermined plans (e.g. robot arm trajectories). The following items of the work plan are intended to extend this research to include learned path planning, which is the next logical level of a control hierarchy.

1. Develop techniques for learned trajectory planning for systems with structural redundancy in the presence of obstacles. Evaluate the robustness of such learned planning systems for incompletely trained models.
2. Develop practical techniques for learning "optimal" trajectories for dynamical systems with constraints. Utility functions such as minimum time, minimum energy, minimum jerk, etc. will be investigated.

Our current work in the area of planning involves using a hierarchical arrangement of CMAC networks to implement a planning model capable of computing relaxed spatial trajectories and adaptively incorporating world imposed constraints. Learning is via direct experience (on-line learning) with the representation of constraints embedded in a hyperspatial representation of a world model mapped directly into the robot's kinematic coordinate system. This hyperspatial representation will largely replace the symbolic world model used in traditional AI systems, which consist of discrete tokens embedded in a database. The trajectories learned are not optimal, but are successful within the constraints imposed. Future work will consider the effect of training paradigms during on-line learning on the nature of the final learned behavior. Preliminary results of this work were presented at the International Joint Conference on Neural Networks in January, 1990 [Rudolph, 1990], a copy of which is included in the appendix.

Reinforcement Learning

We intend to study reinforcement learning within the context of learned biped walking. The immediate goal is to implement a control system which can learn to walk with dynamic balance, for a variety of slopes and payloads, using only crude models of the biped characteristics. The following investigations will proceed using a computer simulation and an experimental biped, both of which have been developed in our laboratory:

1. Develop a reinforcement learning architecture for adapting approximate walking trajectories precomputed using a crude biped model.
2. Evaluate and refine learning technique in the context of efficient biped walking on horizontal surfaces with variable payload.
3. Evaluate and refine learning technique in the context of efficient biped walking on sloped surfaces with variable

payload.

Initial efforts in this section of the research have been aimed at refining both the biped simulator and the physical model. The biped simulator has been updated to more realistically model the effects of foot-strike during walking. The reaction forces involved are important disturbances during actual walking, but are difficult to model accurately due to the related dramatic discontinuity in the system dynamic properties. Other simulators have ignored the transition, formulating independent models appropriate for the foot raised and the foot on the ground, and merely switching from one model to the other. This modeling effort was completed during the report period, and learning control system simulation was begun. The emphasis is on achieving robust walking through learning to modify a predefined walking gait in response to sensory feedback, rather than on trying to learn walking from initial random or otherwise arbitrary movements.

Progress has also been made with the experimental biped, although learning experiments have not yet begun. The focus of efforts during the project period was on developing the necessary support software for the real time control of the biped and data logging proprioceptive and control parameters. Prototype force sensing feet were also developed which provide important feedback concerning the distribution of forces on the sole of the foot, which reflects the state of balance of the system.

Control System Analysis

Our past work has emphasized the use of neural networks as nonlinear models for adaptive control. The concepts have been demonstrated to be practical and effective for difficult simulated and real time experimental control problems. We plan to extend the theoretical analysis of these important learning control techniques in the following areas, using simplified system models:

1. Analyze learning control system performance in the presence of noisy measurements.
2. Develop stability criteria for the closed-loop learning control system in competing control architectures.
3. Analyze convergence times for the neural network weights, and settling times for the performance of the closed-loop learning control system.
4. Compare neural network based learning control with other adaptive control techniques.

In preliminary work performed during the first half of 1989, we developed a protocol for comparing CMAC based neural network control with the "traditional" techniques of model reference adaptive control and the self-tuning regulator. The goal of this work was to contrast the performances of these techniques in closed loop control problems with both linear and nonlinear plants, and with varying degrees of noise. Preliminary results [Kraft and Campagna, 1989a; 1989b; 1990] indicated that the CMAC approach was better in cases with model mismatch and noise, while the traditional techniques converged faster if the plant model was well known and the measurements were low in noise. This simulation work was extended during the current project period, reinforcing the same conclusion. It was realized that many proposed architectures for neural network based control have analogs in the traditional adaptive control literature, using relatively simple parametric models rather than neural network models. Initial efforts were thus made to relate different learning control architectures with different traditional adaptive control architectures. It is hoped that analysis techniques developed for the adaptive control techniques can then be extended to provide insight into the characteristics of the corresponding learning control architectures. This work will continue during the next project period.

Another effort which was started prior to this project but was continued during the project period involved the analysis of the closed loop stability of CMAC neural network based controllers with on-line learning. Closed loop poles have been derived for simple plants and repetitive trajectories [Kraft et al., 1989]. Using such analyses, the effects of parameters such as learning system gain on pole placement can be examined. While the initial analysis was performed using a very simple model, the theoretical results predict features observed in real experiments with more complicated plants. This work is being extended to analyze more interesting models.

Fault Tolerance

Fault tolerance is often mentioned as an attribute of artificial neural networks, although much of the current evidence is anecdotal in nature. Since fault tolerance is clearly an important feature for robust control, we plan to investigate neural network fault tolerance explicitly, as follows:

1. Operational fault tolerance. Study relationships between network size, degree of generalization, function complexity and function reproduction accuracy for networks with faults imposed after training. Using measures of complexity derived from information theory, develop unified bounds to the complexity of networks that will realize a function of given complexity, with a desired approximation accuracy, in the presence of faults.
2. Learning fault tolerance. Study relationships between learning convergence time, final approximation accuracy, and function complexity for faults imposed before training. Using the analogy between learning and system identification, extend known results on the identifiability of input-output systems to situations in which the identification system has parameter constraints (i.e. faults), and apply these to the determination of learning complexity.
3. Fault tolerance enhancement. Develop fault tolerant quantization schemes for a fixed input layer of the network, to create robust internal representations. Design internal representations based on results from diophantine approximation theory to retain representation accuracy in the presence of faults. Study techniques for "weight balancing" to minimize sensitivity to faults.

Initial efforts in this regard have focused on analyzing the operational fault tolerance of CMAC neural networks [Carter et al., 1989]. The general test scenario was one in which weight faults were introduced after network training. Networks were then retrained with the faulty weights. It was found that the sensitivity of the network to such faults was different for "zero weight" versus "saturated weight" faults, and depended on the extent of the network spatial generalization (receptive field size) relative to the dominant spatial wavelengths of the function being learned. Related work demonstrated a previously unknown problem in CMAC weight convergence during training, again related to size of the generalization regions relative to the dominant spatial wavelengths of the function being learned [Carter et al., 1990]. Preliminary results indicate that CMAC networks with tapered receptive fields are less sensitive to such effects than traditional CMAC networks. These issues are actively being pursued.

Related Work in Progress

We recently completed a very high speed implementation of the CMAC neural network using dedicated CMOS logic, rather than a general purpose or RISC processor (ONR grant N00014-89-J-1686). This technology was then used to implement two general purpose CMAC associative memory boards for the industry standard VME bus, facilitating future development of real time applications of neural networks to learning control systems, pattern recognition, and signal processing. Two prototype VME boards have been constructed, each implementing a CMAC network with one million adjustable weights. The boards are currently undergoing exhaustive testing and final support microcode development. VME bus response times for typical CMAC networks with 32 integer inputs and 8 integer outputs are on the order of 200 to 400 microseconds, depending on the network generalization parameter, making the networks sufficiently fast for most robot control problems, and many pattern recognition and signal processing problems. The hardware developed will be evaluated experimentally by the Robotics Laboratory at the University of New Hampshire and by the Robot Systems Division of the National Institute of Standards and Technology. The availability of the high speed hardware will have a positive effect on future experimental research within the scope of this grant. A copy of the most recent progress report for the CMAC hardware project is included in the appendix.

We are currently developing a new experimental testbed which involves two five axis robotic arms, with grippers, in a cooperative working arrangement. The work cell will also include a stereo vision system mounted on a third five axis arm, to provide for active visual sensing during task performance (the cameras can be transported and rotated to achieve the best view). The stereo camera pair will have automatically adjustable parallax angle to provide for robust depth perception. The completed work cell, with two manipulators and

active vision, will provide an excellent environment for testing learning control concepts including path planning, reinforcement learning, active sensing, and distributed/hierarchical control architectures. Issues of hand-eye coordination, multi-arm cooperation, and moving obstacle avoidance can be studied in a system of realistic complexity. Traditional kinematic modeling approaches would be difficult to apply to such a system in which the visual frame of reference is highly variable during task execution. The equipment budget of this grant was utilized to purchase two small CCD video cameras and an 80386 based workstation for this testbed. Two of the robotic arms to be used were purchased from other sources. The third robotic arm and video image processing equipment were already available in the laboratory.

As part of a NSF funded project, we have recently completed a new demonstration of learned hand/eye coordination using our GE P-5 industrial robot. This demonstration involves learning to push a one wheeled vehicle (similar to a chair castor) around a closed track of variable shape using visual feedback. The novel feature of this demonstration is that the vehicle dynamics are unstable (i.e. the front wheeled vehicle wants to rotate around when pushed from behind). In previous demonstrations, on-line learning has been used to improve control system performance for an initially stable system. In this case, the system is unstable without on-line learning and the test track can not be followed even poorly. Rapid on-line learning using CMAC provides the initial stability needed to follow the track at all, and continued practice provides the refined learning required to follow the track accurately.

Bibliography

Albus, J. S., "A New Approach to Manipulator Control: The Cerebellar Model Articulation Controller (CMAC)." Trans. of the ASME, Journal of Dynamic Systems, Measurement and Control, vol. 97, pp. 220-227, September, 1975.

Albus, J. S., "Mechanisms of Planning and Problem Solving in the Brain." Mathematical Biosciences, vol. 45, pp. 247-293, August, 1979.

Jordan, M. I., "Supervised learning and systems with excess degrees of freedom." COINS Technical Report 88-27, University of Massachusetts, 1988.

Carter, M. J., Rudolph, F. J., and Nucci, A. J., "Operational Fault Tolerance of CMAC Networks." Presented at the 1989 IEEE Neural Information Processing Systems Conf., Denver, CO., November, 1989. To appear in Advances in Neural Information Processing Systems 2, San Mateo, CA, Morgan Kaufman, 1990.

Carter, M. J., Nucci, W. T., Miller, W. T., and Rudolph, F. J., "Slow Learning Scenarios for Locally Generalizing Neural Networks and Implications for Fault Tolerance." To be presented at the 24th Ann. Conf. on Information Sciences and Systems, Princeton Univ., March 21-23, 1990.

Hewes, R.P., and Miller, W.T. "Practical demonstration of a learning control system for a five axis industrial robot." Proceedings SPIE: Intelligent Robots and Computer Vision, vol. 1002, 1988.

Kraft, L. G., An, E., and Campagna, D. P., "Comparison of CMAC Controller Weight Update Laws." Proc. 28th Conf. on Decision and Control, Tampa, Fla., December, 1989, pp. 1746-1747.

Kraft, L. G., and Campagna, D. P., "A Comparison of CMAC Neural Network and Traditional Adaptive Control Systems." Proc. of the 1989 American Controls Conf., Pittsburgh, Pa., May, 1989.

Kraft, L. G., and Campagna, D. P., "Comparison of Convergence Properties of CMAC Neural Network and Traditional Adaptive Controllers." Proc. 28th Conf. on Decision and Control, Tampa, Fla., December, 1989, pp. 1744-1745.

Kraft, L. G., and Campagna, D. P., "A Summary Comparison of CMAC Neural Network and Traditional Adaptive Control Systems." To appear in the IEEE Control Systems Magazine, 1990.

Miller, W. T., "A Nonlinear Learning Controller for Robotic Manipulators." Proc. of the SPIE: Intelligent Robots and Computer Vision, vol 726, pp. 416-423, October, 1986.

Miller, W. T., "A Learning Controller for Nonrepetitive Robotic Operations." Proc. of the Workshop on Space

Telerobotics, JPL publication 87-13, vol. II, pp. 273-281, Pasadena, CA, January 19-22, 1987.

Miller, W. T., "Sensor Based Control of Robotic Manipulators Using A General Learning Algorithm." IEEE J. of Robotics and Automation, vol. RA-3, pp. 157-165, April, 1987.

Miller, W. T., Glanz, F. H., and Kraft, L. G., "Application of a General Learning Algorithm to the Control of Robotic Manipulators." The International Journal of Robotics Research, vol. 6.2, pp. 84-98, Summer, 1987.

Miller, W.T., and Hewes, R.P.: Real time experiments in neural network based learning control during high speed, nonrepetitive robot operations. Proceedings of the Third IEEE International Symposium on Intelligent Control, Washington, D.C., August 24-26, 1988.

Miller, W. T., "Real time learned sensor processing and motor control for a robot with vision." Proceedings of the First Annual Conference of the International Neural Network Society, Boston, MA, September 6-10, 1988, pp. 347.

Miller, W. T., "Real Time Application of Neural Networks for Sensor-Based Control of Robots with Vision." IEEE Transactions on Systems, Man, and Cybernetics special issue on Information Technology for Sensory-Based Robot Manipulators, Vol. 19, pp. 825-831, 1989.

Miller, W. T., Hewes, R. P., Glanz, F. H., and Kraft, L. G., "Real Time Dynamic Control of an Industrial Manipulator Using a Neural Network Based Learning Controller." To appear in the IEEE J. of Robotics and Automation, 1990).

Rudolph, F., "Locally Optimizing Neural Networks in Adaptive Path Planning." Presented at the IJCNN, Washington, D.C., Jan. 15-19, 1990.

Werbos, P. J., "Neural Networks for Control and System Identification." Proc. 28th Conf. on Decision and Control, Tampa, Fla., December, 1989, pp. 260-265.